

PCT


WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G11B 5/09, 3/90	A1	(11) International Publication Number: WO 98/03970 (43) International Publication Date: 29 January 1998 (29.01.98)
(21) International Application Number: PCT/US97/12503 (22) International Filing Date: 18 July 1997 (18.07.97) (30) Priority Data: 08/684,348 18 July 1996 (18.07.96) US (71) Applicant: QUANTUM CORPORATION [US/US]; 500 McCarthy Boulevard, Milpitas, CA 95035 (US). (72) Inventors: ROBERTSON, Bryan; 6886 Chantel Court, San Jose, CA 95129 (US). HOSPODOR, Andrew, D.; 208 Kimble Avenue, Los Gatos, CA 95032 (US). (74) Agents: CHUN, Debra, A. et al.; Quantum Corporation, 500 McCarthy Boulevard, Milpitas, CA 95035 (US).		(81) Designated States: AU, CA, CN, JP, KR, SG, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published With international search report.

(54) Title: HARD DISK DRIVE HAVING DYNAMIC IN-LINE SPARING



(57) Abstract

In a hard disk drive (100) including a disk spindle (102) and a rotary actuator assembly (104), a method for compensating for a defective sector in a track (107-110) is provided. During a read operation whereby data is retrieved in a steady stream (e.g., an audio/video drive), a quick check is made each time a sector is accessed to determine whether the data is correct. If a sector is deemed to be defective, the disk drive sets an allocation flag and records it in a diagnostic memory (111). The disk drive continues with its read operation in this manner until all requested data has been retrieved. Thereupon, the disk drive goes back and performs a reallocation procedure to correct for any defective sectors as indicated by the allocation flags retained in the diagnostic memory. The described method preserves the sequential nature of data on the disk (101), and provides assurance that the reallocation process continues without data loss, even during power sequencing.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

HARD DISK DRIVE HAVING DYNAMIC IN-LINE SPARING

FIELD OF THE INVENTION

5

The present invention pertains to hard disk drive 100s having the capability of performing dynamic in-line sparing.

BACKGROUND OF THE INVENTION

10

Hard disk drive 100s offer a convenient, efficient, and cost effective way to store vast amounts of data. Hence, hard disk drive 100s are found in virtually every computer system, ranging from mainframes, workstations, and desktop personal computers to portable laptop computers. Basically, in a hard disk drive 100, data is stored magnetically as a series of digital bits consisting of "1's" and "0's." A transducer or "head" is used to record the data onto a magnetically coated disk, whereby a "1" is represented by one magnetic polarity and a "0" is represented by an opposite polarity. Later, the same head is used to read pre-recorded data back from the disk. Often, several of these disks are stacked together and rotated about a spindle by a spindle motor. As the transducer is moved radially by a servo/actuator assembly across the surface of a spinning disk, a number of circular "tracks" are defined. These tracks are subdivided into a number of "sectors." It is within these sectors that the digital bits of data are stored by controlling the polarities of magnetic domains contained in the sectors.

One problem frequently encountered with hard disk drive 100s relates to that of data corruption. Occasionally, some or all of the original data is unintentionally altered or lost. For instance, small portions of the magnetic disk might lose their magnetic cohesion over a period of time. These grown defects might cause data in one or more sectors to become unreliable. In the past, this problem was fixed by placing defective sectors off limits. In other words, internal circuitry and firmware routines were used to detect certain sectors which

have been corrupted or compromised. The problem sectors would then be isolated thus preventing data from being written to and read from those sectors. Some hard disk drive 100 manufacturers have even taken the step of implementing error correction code to correct the bad data and reallocating the corrected data to a different sector which is functioning properly.

However, this prior art reallocation process cannot be practiced upon disk drive 100s dedicated to storing audio or video (A/V) data. A/V drive 100s are finding great popularity as a means for storing audio and video information, especially relating to multimedia applications. The reason why the prior art reallocation process cannot be readily applied to A/V drive 100s is due to the real-time nature of audio and video. Audio as well as video data are magnetically stored within the hard disk drive 100 as a continuous bit stream. This data is then subsequently read back from the disk drive 100 for playback at the discretion of the user. Any significant interruption in the readback of the data would result in perceptible audio glitches and visually disturbing video artifacts. In a typical prior art reallocation process, approximately 66 msec in rotational latency is spent retrying a sector six times; mapping the defect out requires another rotation to mark the sector (4.17 msec) and a seek (8.5 msec) to write the data to the new sector; and additional time is incurred for command overhead (20 μ sec) and data processing (34 μ sec). Hence, a total time of 79.21 msec is spent in performing a typical prior art reallocation sequence. This momentary interruption in the supply of data from the hard disk drive 100 to the computer system will produce audible clicks and pops as well as annoying visual discrepancies. Thus, A/V disk drive 100 manufacturers and users are faced with a no-win situation. If bad sectors are not reallocated, the problem grows worse as more and more sectors become defective over time. However, fixing bad sectors also introduces unwanted glitches and video artifacts.

Thus, there is a need in the prior art for some mechanism that compensates for bad sectors without introducing any latencies into the retrieval of audio/video data from the disk

drive 100. The present invention provides one solution to the problem by correcting defects off-line to improve data integrity and overall drive 100 performance, whereby read operations are completed without any interruptions.

5 SUMMARY OF THE INVENTION

The present invention pertains to a hard disk drive 100 having the capability of essentially eliminating defective sectors. During a read operation whereby data is retrieved in a steady stream (e.g., an audio/video drive 100), a quick ECC check is made
10 each time a sector is accessed to determine whether the data is correct. If data from a particular sector is consistently incorrect, that particular sector is deemed to be defective. Thereupon, the disk drive 100 stores the location of the defective sector into a diagnostic memory which is periodically written on a disk surface. This can be accomplished rather quickly with negligible impact to the read operation. The disk drive 100 then continues
15 with its normal read operation. Each time a defective sector is encountered, the location of the defective sector is recorded into the diagnostic memory. The diagnostic memory is periodically written onto a reserved storage space on one of the disk storage surfaces. The defective sector recording process is repeated until all of the requested data have been retrieved. Thereupon, the disk drive 100 systematically performs a reallocation procedure
20 on all tracks containing a defective sector as indicated by the diagnostic memory.

In one embodiment of the present invention, the reallocation process consists of copying all of the data stored in the track containing a defective sector into a specially reserved diagnostic track or tracks. In-line sparing is then performed on the track containing
25 the defective sector. The in-line sparing process reformats the logical sectors of that track so that data is prevented from ever being written to or read from the defective sector. The disk drive 100 will now skip the defective sector and go straight to the next sector. Lastly, the data is read from the diagnostic track and written back to the reformatted track. The disk

drive 100 repeats this process until all tracks containing defective sectors have been successfully reallocated. In the present invention, care is taken to ensure that a copy of the data is always maintained. Furthermore, in the event of an unexpected disruption, the disk drive 100 restarts the reallocation process at an appropriate point so that data integrity is preserved.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

Figures 1A (plan view) and 1B (side view) illustrate diagrammatically some of the elements of a generic hard disk drive within which the present invention may be practiced.

Figure 2 shows a flowchart describing the steps for initiating a reallocation process.

Figure 3 shows how a track with a defective sector is reallocated.

Figure 4 is a flowchart describing the steps for performing a reallocation process according to the currently preferred embodiment of the present invention.

Figure 5 shows a flowchart describing the steps for recovering from any disruptions occurring during the reallocation process.

DETAILED DESCRIPTION

A hard disk drive having dynamic in-line sparing is described. In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be obvious, however, to one skilled in the art that the present invention may be practiced without these specific
5 details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid obscuring the present invention. Furthermore, although the present invention is ideally suited for audio/visual (A/V) hard disk drives, the present invention may also be applied in other types of drive 100s (e.g., generic hard disk drives, magneto-optic disk drive 100s, optical disk drives, floppy disk drive 100s, etc.)

10

Figure 1A shows the top view of a generic hard disk drive 100 within which the present invention may be practiced. As shown in the Figure 1B side view in elevation, a number of magnetically coated disks 101 are stacked together and rotated about a spindle 102 by an in-hub spindle motor (hidden from view). A number of arms 103 of e.g. a rotary voice
15 coil actuator assembly 104 are used to move read/write transducers 105 radially across the data storage surfaces of the disks 101. While six disks are shown in Figure 1B, greater or fewer disks may be employed.

The disk drive 100 additionally includes circuitry 106 implementing a servo system
20 to control the movement of actuator assembly 104 for performing seeks in read/write operations. The circuitry 106 also provides signal processing, decoding, and interfacing functions. Furthermore, the circuitry 106 includes either a microcontroller or microprocessor which performs the reallocation and dynamic in-line sparing process of the present invention. The disk drive 100 circuitry 106 is connected via a conventional bus structure 115 to a host
25 computing system 116. The disk spindle 102 and rotary actuator assembly 104 are secured to a base plate 113, and a cover 114 isolates an interior space from ambient external conditions to provide a clean operating environment.

Numerous concentric data storage tracks are defined on surfaces of each of the disks 101. A single disk may contain thousands of such tracks on each storage surface. Four exemplary tracks are shown as 107-110 in the Figure 1A plan view. These tracks are further formatted into a plurality of sectors. It is within these sectors in which data is ultimately stored. The data for audio and video content is stored in a continuous bit stream. Often, such data occupies several tracks. The data is written to these tracks in a sequential fashion so that the data may be retrieved with a minimal number of seeks. By writing the data sequentially, the computer system may retrieve the data in a smooth, uninterrupted flow that is mandatory for high quality video display and audio playback.

10 In the present invention, the circuitry 106 is used to detect whether a sector has become defective. A quick determination is made at the time of a read operation to determine whether the data is correct. If the data is incorrect, then firmware embedded in the circuitry 106 may optionally attempt to fix the data by invoking an on-the-fly error correction code (ECC) process. If the error still exists, the disk drive 100 may then attempt a multiple error burst ECC process. If it is determined that the error persists, then the present invention marks the defective sector and continues with its current read operation. Only after the complete read operation is finished does the disk drive 100 then go back and reallocate the defective sector. In a reallocation process, the original track containing the defective sector is read and temporarily stored onto a diagnostic track. An in-line sparing operation is performed on the original track. The in-line sparing operation renders the defective sector transparent to the disk drive 100 so that it is skipped. This is accomplished by disabling the defective sector and utilizing a spare sector at the end of the track. The original track is then reformatted so that the sectors are logically sequential. The data is then read back from the diagnostic track and written into the reallocated track. Thereby, the disk drive 100 automatically skips over the defective sector whenever it performs a subsequent read or write operation to that particular track.

As an example, suppose that the host computer system 116 commands the hard disk drive 100 to store a five minute video segment. The hard disk drive 100 finds and allocates four available, consecutive tracks 107-110 and proceeds to write the digitized video information onto those tracks in a continuous bit stream (e.g.,
5 011101001101000011101011100. . .). Later, suppose that the second sector of track 108 eventually becomes defective. The next time that track 108 is read, the defect is discovered. The disk drive 100 takes note of this fact and continues with its read operation until all five minutes worth of video data is read and displayed by the host computer system. Thereupon, the entire track 108 is read and stored onto a diagnostic track 111. While Figure 1A places
10 the diagnostic track 111 at an outer radius of the disk 101, any suitable location, including a track near an inside diameter of the disk, may be used for the diagnostic track (or diagnostic zone of tracks). The circuitry 106 corrects the defective data. An in-line sparing operation is performed on track 108. The data is then read from the diagnostic track 111 and written to the reformatted track 108. Finally, the disk drive 100 notes that in-line sparing has been
15 successfully performed on the defective sector.

Figure 2 shows a flowchart describing the steps for initiating a reallocation process. Initially, a read operation is requested by the host computer system for the retrieval of data from the hard disk drive 100. A seek is performed to place the selected transducer 105 over
20 the first sector containing the desired data. That sector is then read, step 201. A quick determination (e.g., parity, Hamming code, etc.) is made as to whether the data in that sector is correct, step 202. If the data read in step 201 is correct, then step 203 is skipped. Otherwise, if the data is incorrect, an allocation flag is written to a designated diagnostic area on one of the disks, step 203. The allocation flag identifies the location of the defective
25 sector. In step 204, a determination is made as to whether all of the data has been retrieved. If the end of the data has not been reached, step 201 is repeated by reading the next sector. After the entire requested data has been retrieved, the disk drive 100 performs a reallocation process to fix the sectors which were marked as being defective in step 203.

Figure 3 shows how a track with a defective sector is reallocated. In the present invention, each of the tracks (or zones of radially adjacent tracks) with the capability of being reallocated, includes one or more spare sectors. For instance, track 301 is comprised of 187
5 separate sectors and one spare (S) sector. In an A/V drive which may be similar to the disk drive 100 shown in Figures 1A and 1B, data is written sequentially to each successive logical sector (e.g., first, logical sector 1 is written, then logical sector 2, followed by logical sector 3, etc.). If one of the sectors becomes defective, that particular track is reallocated in place the next time it is read. For instance, suppose that physical sector 2 of track 302(a) happens to
10 become defective. The A/V hard disk drive detects this condition and flags sector 2 of track 302(a) as being defective. After the read operation has completed, the A/V disk drive copies the data contained in all of the 187 sectors of track 302(a) onto a diagnostic, or reserved, track 303. In the preferred embodiment, the outer most diameter (OD) track is reserved as the diagnostic track 303. However, those skilled in the art would recognize that another track
15 could be reserved as the diagnostic track 303. The A/V disk drive corrects the defective data and substitutes the corrected data in place of the erroneous data in logical sector 2. Hence, diagnostic track 303 now contains the correct sequence of data. Diagnostic track 303, or a separate track 30x may include storage locations for allocation flags, and any other flags needed in a particular implementation of the present invention.

20

In the meantime, dynamic in-line sparing is performed on the original track 302(b). In-line sparing reformats the original track 302(b) so as to bypass the defective sector 2. This is accomplished by assigning logical sectors to physical sectors, whereby the defective sector is essentially masked out. The sector numbers on top of a track represent those sectors
25 actually present and which physically exist on that track. Physical sectors do not change. The sector numbers within a track represent the logical sectors that are recognized by the disk drive 100. Normally, the logical sector number matches its corresponding physical sector number. However, the logical sector number of a defective physical sector may be realigned

so as to skip the defective physical sector. For example, the result after in-line sparing has been applied to track 302(a) is shown as track 302(b). Track 302(b) shows the distinction between logical and physical sectors after in-line sparing has been performed to mask out defective sector 2. It can be seen that logical sector 1 still corresponds to physical sector 1.

5 However, logical sector 2 now corresponds to physical sector 3. A logical "FF" is written to the defective sector to identify the sector as an "illegal" sector. Each successive logical sector is shifted over by one. Hence, logical sector 2 corresponds to physical sector 3; logical sector 3 corresponds to physical sector 4; etc. The physical spare sector (S) now becomes logical sector 187. Since the A/V disk drive performs read/write operations on a logical sector basis,
10 it will skip over the defective physical sector. For instance, in a write operation, data will be written sequentially to logical sectors 1-187. This means that logical sector 1 causes data to be written to physical sector 1. Logical sector 2 causes data to be written to physical sector 3; logical sector 3 causes data to be written to physical sector 4; etc. Thereby, the defective physical sector 2 is effectively skipped. Likewise, in a read operation, the data will be read
15 according to logical sectors 1-187. This means that the data will be read from physical sector 1, then from physical sector 3, and physical sector 4, etc. Again, the defective physical sector 2 is skipped.

While the term "track" is used herein to denote a single circular data storage locus on
20 a storage surface, those skilled in the art will appreciate that a track locus could be a single spiral across the entire data storage surface. Alternatively, tracks may be arranged into zones, with each zone having a data transfer rate optimized to disk radius (assuming a constant disk rotational velocity). Thus, the term "track" as used herein should be understood to include a
"zone" arrangement.

25

Figure 4 is a flowchart describing the steps for performing a reallocation process according to a preferred embodiment of the present invention. Once the current read process is finished and the disk drive is idle for four times the anticipated reallocation period, the off-

line reallocation process may proceed. In performing this process, the reallocation combines an off-line scheduling feature with live in-line sparing. As discussed above, allocation flags are written to disk in the diagnostic area during a read operation to flag those sectors which were determined to be defective. In addition, a power on (POR) flag is also stored within the
5 diagnostic area. The POR flag is used to indicate when the drive loses power during the reallocation process. The allocation and POR flags may share a byte.

In the first step 401, the original track containing the defective sector as indicated by an allocation flag stored in diagnostic memory, is read. An attempt is made to correct the
10 defective data, if possible. Next, the data read from the original track is written to the diagnostic track, step 402. In one embodiment, the diagnostic track contains an $OD_{\text{maxsectors}+1}$ sectors and may occupy one or more physical tracks of the disk drive. The diagnostic track is then verified against the original data track, step 403. Furthermore, the POR flag is set to an ON condition in order to write to the disk, step 404. Finally, the data is
15 written to the disk, step 405. In step 406, a determination is made as to whether the data track is re-written. If it is re-written, then the process begins anew at step 401. Otherwise, the process proceeds to step 407.

In step 407, a mutex lock is applied to the original data track. The mutex lock
20 prevents access to that particular track. In other words, the disk drive cannot perform read or writes to that problem track when the mutex lock is applied. Thereupon, the sector identifications (IDs) on the original track are reformatted, step 408. This reformatting step is also referred to as in-line sparing. In one embodiment, dynamic in-line sparing involves using that sector flag on a recurring basis. This flag contains either the cylinder/head/sector
25 (CHS) or logical block address (LBA) for the defective sector. For drives with ID-less addressing, the wedge count is used. This flag is queued, by storing the flag in a buffer and not performing retries or triple burst ECC, for later retries (after the current read request is complete). In step 409, the data is read back from the diagnostic track. This data is then

written onto the reformatted original track. step 410. The data is verified in step 411 to check that it was written properly. Next, the allocation and POR flags are cleared from memory, step 412. Finally, the mutex lock is removed, step 413. If there is another defective sector, as determined in step 414, the reallocation process is repeated at step 401 for the track
5 containing the defective sector. Steps 401-413 are repeated until all defective sectors have been reallocated or until there are no more spare sectors. It should be noted that data is never jeopardized by the present invention. Care has been taken to guarantee that a copy of the data is always retained someplace on the disk at all times.

10 Figure 5 shows a flowchart describing the steps for recovering from any disruptions occurring during the reallocation process. This makes the present invention fail-safe in the event of a power cycle. In step 501, the disk drive monitors all disruptions. Examples of disruptions include power cycles (e.g., power-on conditions, power-off conditions, disk crashes, etc.). If a disruption occurs during the reallocation process, circuitry within the disk
15 drive determines the exact point of the reallocation process at which the disruption occurred, step 502. Referring back to Figure 4, if the disruption occurs during steps 401-406, then the disk drive restarts at step 401 at the next possible opportunity that arises. For example, if the disk drive is accidentally powered down while it is in the middle of performing steps 401-406, then the disk drive executes step 401 the next time that the disk drive is powered back
20 on. This situation is shown as steps 503 and 504. Steps 505-506 show that if the disruption occurs in the middle of steps 407-411, then the reallocation process is restarted at step 407 at the next possible opportunity. Otherwise, if the disruption occurs during steps 412-413, then the reallocation process is restarted at step 412 at the next earliest opportunity (see steps 507-508). Thus, the present invention guarantees that data integrity is strictly maintained.

25

The foregoing descriptions of specific embodiments of the present invention have been presented for purposes of illustration and description. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed, and obviously many

modifications and variations are possible in light of the above teaching. The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, to thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the claims
5 appended hereto and their equivalents.

What is claimed is:

1. In a disk drive storage system having at least one disk with a plurality of concentric data tracks, each track being formatted into a plurality of sectors for storing data, a method for compensating for a defective sector, the method comprising steps of:
 - (a) during a read operation including multiple accesses to sequential sectors:
 - (i) reading data from a first sector of the track,
 - (ii) determining whether the read sector is defective,
 - (iii) storing an indication that the read sector is defective into a diagnostic memory if the read sector is defective,
 - (iv) determining whether all data for the read operation has been accessed,
 - (v) if all data has not been accessed, reading data from a next sector of the track and repeating steps (ii) - (iv) until all data has been accessed;
 - (b) retrieving from the diagnostic memory indications that the read sectors are defective; and
 - (c) reallocating the data contained in the defective sectors, wherein sequential nature of data is preserved.
2. The method of claim 1, wherein the reallocating step comprises steps of:
 - reading all data from the track;
 - writing the data from the track onto a reserved track;
 - performing an in-line sparing operation on the track to reformat the track;
 - reading the data back from the diagnostic track; and
 - writing the data onto the reformatted track.
3. The method of claim 2, wherein the reallocating step further comprises a step of performing error correction to correct data corresponding to the defective sector.

4. The method of claim 2, wherein the reallocating step further comprises steps of:

- preventing access to the data during reallocation;
- setting a power-on flag;
- 5 verifying that the data was properly written to the reformatted track;
- clearing the power-on flag;
- enabling access to the track.

5. The method of claim 2 further comprising steps of:

10 determining whether a disruption occurs during the reallocating step;

restarting the reallocating step at different points according to when the disruption occurred, wherein data integrity is maintained.

6. The method of claim 1, wherein the indication that the read sector is

15 defective is represented by an allocation flag.

7. The method of claim 6 further comprising the step of storing a plurality of queued allocation flags in the diagnostic memory, wherein the allocation flags indicate which sectors of the disk are currently defective.

20

8. In a data storage system having a magnetic medium formatted into a plurality of sectors for storing data, a method for correcting for a defective sector of the magnetic medium, the method comprising steps of:

(a) performing a read operation, wherein data is read from a plurality of sectors of

25 the magnetic medium and during the read operation:

- determining whether data read from one of the sectors is defective,
- wherein if one of the sectors is defective, location of the defective sector is stored in a memory;

- (b) locating the defective sector;
- (c) writing the data from the defective sector to a designated sector of the magnetic medium;
- (d) reformatting the defective sector to prevent data from being written to or read
5 from the area that is defective; and
- (e) writing the data from the designated sector of the magnetic medium to the reformatted sector.

9. The method of claim 8, wherein the reformatting step comprises a step of
10 performing an in-line sparing operation.

10. The method of claim 8, wherein the reformatting step further comprises a
step of performing error correction to correct data in the defective sector.

15 11. The method of claim 8, wherein the location corresponds to an address
specifying cylinder, head, sector information.

12. The method of claim 8, wherein the location of the defective sector
corresponds to a logical block address.

20

13. The method of claim 8, wherein the location of the defective sector
corresponds to a wedge count.

14. The method of claim 8, wherein the reformatting step further comprises
25 steps of:

- preventing access to the defective sector during the reformatting process;
- setting a power-on flag;
- verifying that the data was properly written to the reformatted sector;

clearing the power-on flag; and
allowing access to the reformatted sector.

15 15. The method of claim 8 further comprising steps of:
determining whether a disruption occurs during the reformatting step;
restarting the reformatting step at different points according to when the disruption
occurred, wherein data integrity is maintained.

10 16. The method of claim 8 further comprising a step of storing a plurality of
queued allocation flags in memory, wherein the allocation flags indicate which sectors of
the disk are defective.

15 17. In a hard disk drive in which data is stored in sectors of a track defined on
a magnetic disk, a method of assuring data integrity for a defective sector of a track, the
method comprising steps of:

20 a) reading data corresponding to the track containing the defective sector;
b) writing the data to a diagnostic track on the disk;
c) preventing access to the track containing the defective sector;
d) reformatting the track containing the defective sector to produce a reformatted
track;

25 e) reading the data from the diagnostic track;
f) writing the data to the reformatted track; and
g) allowing access to the reformatted track,
wherein the defective sector is reallocated and data is prevented from being
written to or read from the track containing the defective sector while the defective sector
is being reallocated, thereby assuring data integrity.

18. The method of claim 17 further comprising steps of:

restarting the method at step a if a disruption occurs during steps a - b;
restarting the method at step c if a disruption occurs during steps c - f; and
restarting the method at step g if a disruption occurs during step g.

- 5 19. The method of claim 17 further comprising steps of:
 verifying that the data was written properly to the diagnostic track in-between
 steps b and c;
 verifying that the data was written properly to the reformatted track in-between
 steps f and g.

10

20. A hard disk drive comprising:
 at least one magnetic disk including a diagnostic memory area, the disk having a
 plurality of tracks formatted into a plurality of sectors for storing digital data, the
 diagnostic memory area for storing a location of a defective sector;
15 a transducer for reading and writing the digital data to the sectors of the magnetic
 disk;
 an actuator assembly coupled to the transducer for positioning the transducer over
 a designated track;
 a controller for detecting defective sectors read during a read command, and for
20 reallocating the defective sectors after all sectors corresponding to the read command
 have been read.



25

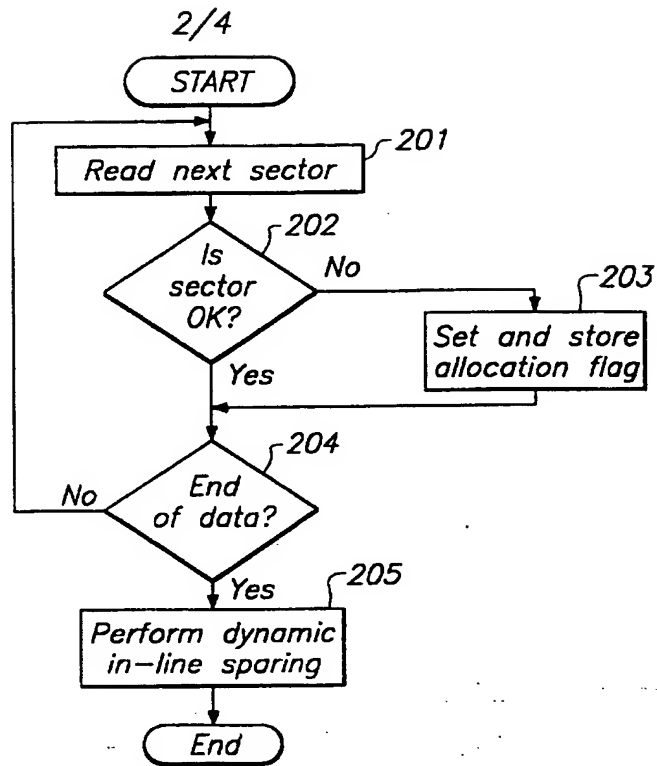
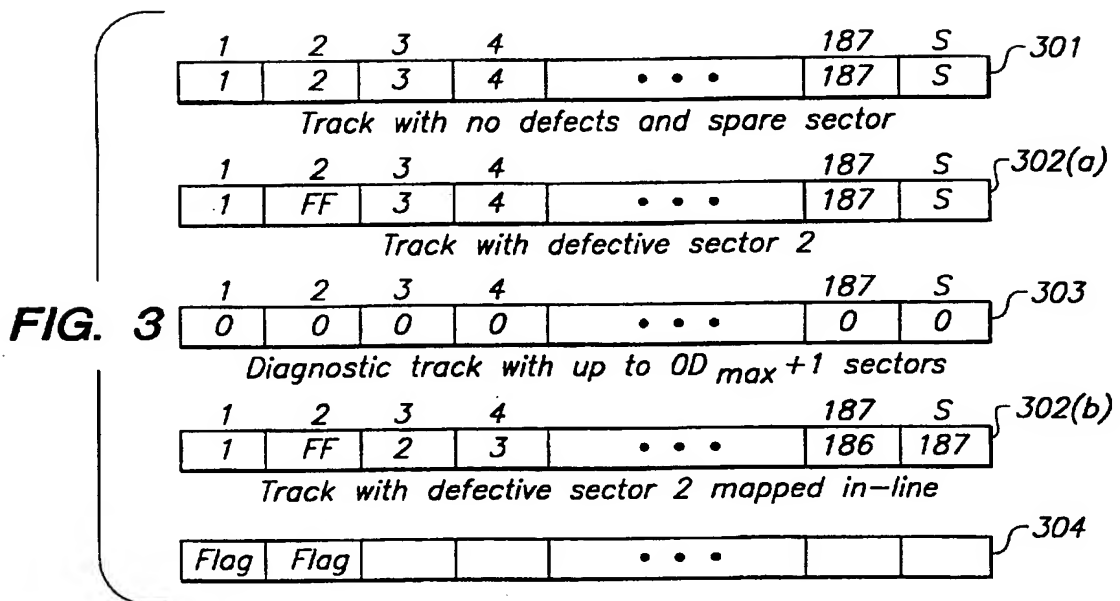
21. The hard disk drive of claim 20, wherein locations of defective sectors of
 the disk are represented by sector flags which are queued in the diagnostic memory area.

22. The hard disk drive of claim 20, wherein the controller performs in-line
 sparing operations to reformat each track containing defective sectors according to the
 sector flags stored in the diagnostic memory area.

23. The hard disk drive of claim 20, wherein if the reallocation process is interrupted, the controller notes a point during the reallocation process at which the interruption occurs and restarts the reallocation at a next available opportunity so that data
- 5 integrity is maintained.

1/4

**FIG. 1A****FIG. 1B**

**FIG. 2**

3/4

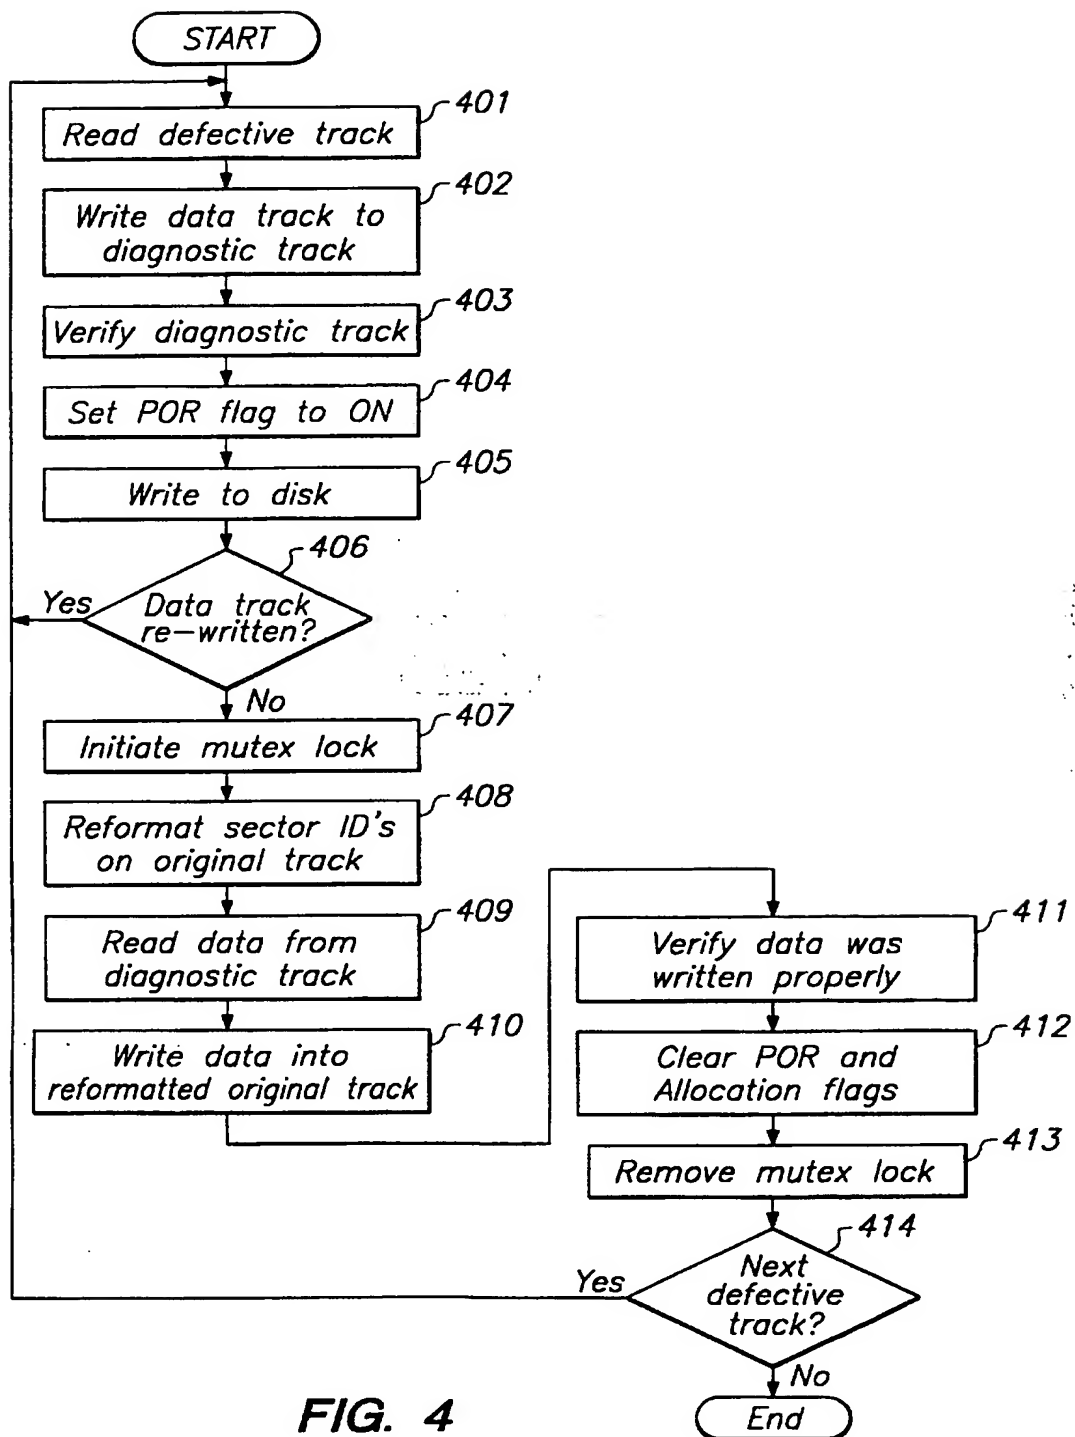
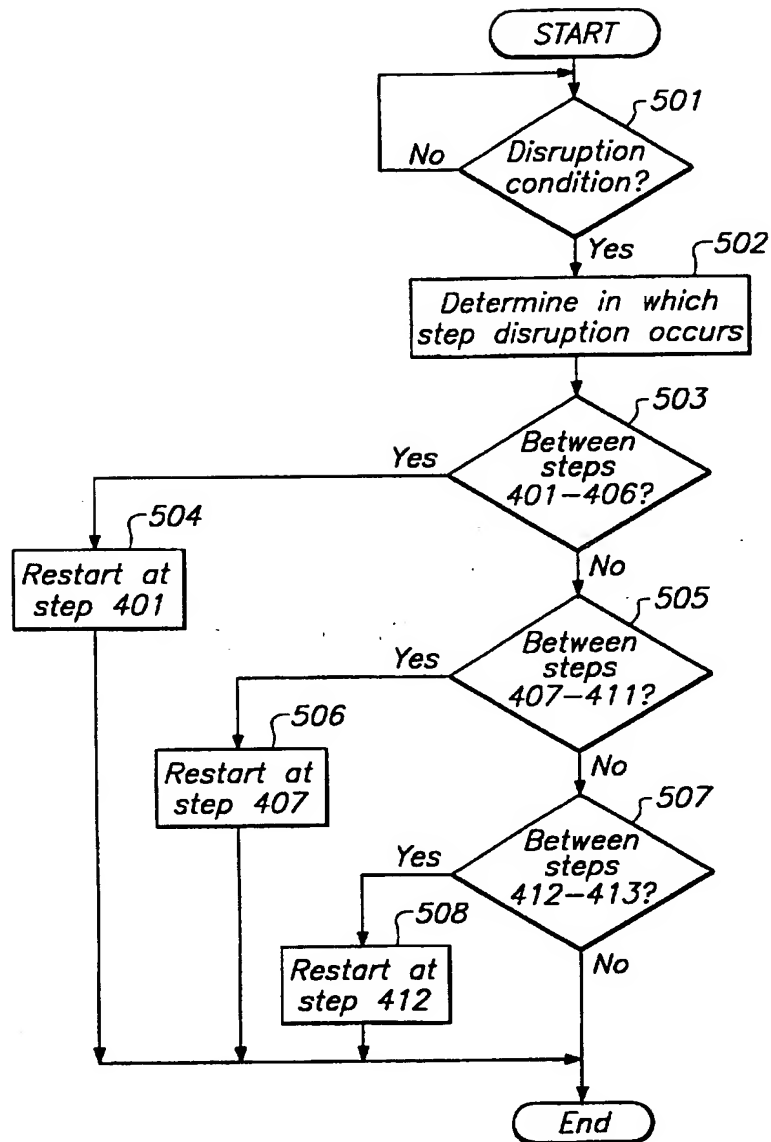


FIG. 4

4/4

**FIG. 5**

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US97/12503

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G11B 5/09, 3/90

US CL : 360/53; 369/54, 58

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 360/53; 369/54, 58

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5,237,553 A (FUKUSHIMA ET AL) 17 August 1993 (17/08/93), see the entire document especially figs. 5-7 and col. 9 lines 39-68.	1-23
X	US 5,166,936 A (EWERT et al) 24 November 1992 (24/11/92), see figs. 1-10.	1-11, 14-23
Y	US 5,523,903 A (HETZLER ET AL) 04 June 1996 (4/6/96), see col. 8 line 45 - col. 9 line 12.	12-13

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

* Special categories of cited documents:

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document published on or after the international filing date

L document which may throw doubt on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O documents referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T

later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X

document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y

document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

Z

document member of the same patent family

Date of the actual completion of the international search

12 SEPTEMBER 1997

Date of mailing of the international search report

17 OCT 1997

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

REGINA NEAL

Telephone No. (703) 308-4078

THIS PAGE BLANK (USPTO)